

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
28 November 2002 (28.11.2002)

PCT

(10) International Publication Number  
**WO 02/096012 A1**

(51) International Patent Classification<sup>7</sup>: **H04L 9/00**

(21) International Application Number: PCT/US02/15163

(22) International Filing Date: 13 May 2002 (13.05.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
09/861,264 18 May 2001 (18.05.2001) US

(71) Applicant: **CLAYMORE SYSTEMS, INC.** [US/US];  
3246 Louis Road, Palo Alto, CA 94303 (US).

(72) Inventors: **DICK, Kevin, Stewart**; 3246 Louis Road, Palo  
Alto, CA 94303 (US). **RESCORLA, Eric, Kenneth**; 2439  
Alvin Street, Mountain View, CA 94043 (US).

(74) Agents: **BHAKAR, Vidya, R.** et al.; Squire, Sanders &  
Dempsey, 600 Hansen Way, Palo Alto, CA 94304-1043  
(US).

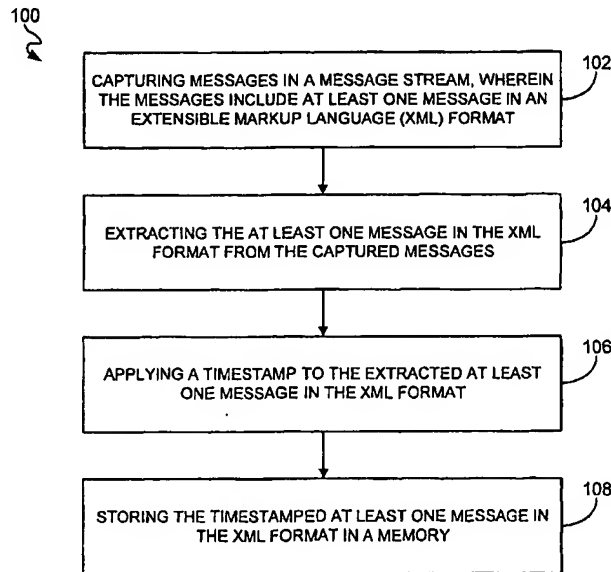
(81) Designated States (*national*): AE, AG, AL, AM, AT, AU,  
AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU,  
CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH,  
GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,  
LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW,  
MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI,  
SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA,  
ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,  
KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW),  
Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),  
European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR,  
GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent  
(BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR,  
NE, SN, TD, TG).

Published:  
— with international search report

For two-letter codes and other abbreviations, refer to the "Guid-  
ance Notes on Codes and Abbreviations" appearing at the begin-  
ning of each regular issue of the PCT Gazette.

(54) Title: SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR AUDITING XML MESSAGES IN A NET-  
WORK-BASED MESSAGE STREAM



(57) Abstract: A system, method and computer program product for auditing a message in a message stream are disclosed. Mes-  
sages in a message stream are captured including at least one message in an extensible markup language (XML) format (102). Each  
message in the XML format is then extracted (104) from the captured messages and has a timestamp applied thereto. Each times-  
tamped message in the XML format is then stored in a memory (108).

WO 02/096012 A1

**SYSTEM, METHOD AND COMPUTER PROGRAM PRODUCT FOR  
AUDITING XML MESSAGES IN A NETWORK-BASED MESSAGE  
STREAM**

5 **FIELD OF THE INVENTION**

This invention relates to document security, and more particularly, relates to verification and authentication of electronic messages.

10 **BACKGROUND OF THE INVENTION**

As enterprises increasingly move their operations from the paper world to the electronic one, they lose critical capabilities of the old paper-based infrastructure. With paper records generated at each stage in a transaction, there was a natural  
15 record of the events. This paper trail provided a number of critical benefits including proof that the transaction occurred as specified at a particular time, an automatic backup of documents at each stage of the transaction, and the ability to file documents in the most appropriate manner or even duplicate them and place them in more than one file at once.

20 By contrast, when enterprises carry out transactions electronically, they lose all these benefits due to the following factors. First, because electronic documents are easily and undetectably modified, it's far easier to tamper with audit logs. Second, employees often incorrectly modify transactions and save them. Once  
25 this occurs, it is often extremely difficult to recover the original transaction, leading to accounting irregularities. Third, documents are stored by whatever program created them in whatever format that program uses.

While the first wave of computerization of business process removed the benefits  
30 of a paper infrastructure, the rise of XML-based Internet business processes allows enterprise to reclaim them. Once transactions occur over the Internet, it becomes possible to capture them in a separate device that then provides long

term secure verification of and access to the message content. The barriers to performing this kind of capture and analysis include:

- Detecting XML messages and extracting them from the surrounding network traffic.
- 5 • Extracting the XML data from the underlying transport.
- Providing high enough throughput under high message loads.
- Maintaining a tamperproof log of all data recorded.
- Accurately determining the time of processing in order to provide reliable auditing.
- 10 •

## SUMMARY OF THE INVENTION

- A system, method and computer program product for auditing a message in a message stream are disclosed. Messages in a message stream are captured
- 15 including at least one message in an extensible markup language (XML) format. Each message in the XML format is then extracted from the captured messages and has a timestamp applied thereto. Each timestamped message in the XML format is then stored in a memory.
- 20 In one aspect of the present invention, the message stream may include a plurality of messages using a variety of protocols. In another aspect, the timestamp may include a digital signature. In a further aspect, the memory may comprise a optical storage medium or a write once storage medium. In an additional aspect, the timestamped message in the XML format may be encrypted prior to storage in
- 25 the memory.

- In one embodiment of the present invention, the captured messages may be parsed to identify each message in the XML format for extraction. In an another embodiment, the message stream may be carried out over a communication path
- 30 having one or more segments and where messages are captured at each segment. In such an embodiment, the captured messages may then be transmitted from each segment to an aggregation module prior to extraction of the messages in the XML

format. In a further embodiment, a report relating to the captured messages may be generated. In an additional embodiment, the message stream may transverse a security boundary having first and second sides where messages on the first side of the security boundary are in an encrypted format and messages on the second  
5 side of the security boundary are in an encrypted format. In this embodiment, an encrypted version of each message in the XML format may be captured, extracted, and timestamped on the first side of the security boundary while a plaintext version of each message in the XML format is captured, extracted, and timestamped on the second side of the security boundary. As a further option, the  
10 encrypted and plaintext version of each message in the XML format may then be correlated to detect any changes between the versions of the respective message.

### BRIEF DESCRIPTION OF THE DRAWINGS

15 Figure 1 is a flowchart of a process for auditing an extensible markup language (XML) message in a message stream in accordance with an embodiment of the present invention;

Figure 2 is a schematic diagram of a system for carrying out a process for auditing  
20 an XML message in a network-based message stream in accordance with an embodiment of the present invention;

Figure 3 is a schematic representation of system for auditing XML messages in a switched network in accordance with an embodiment of the present invention.

25 Figure 4 is a schematic representation of components of a parsing module in accordance with an exemplary embodiment of the present invention;

Figure 5 is a schematic representation of an embodiment of a system capable of  
30 carrying out multiple correlated capture in accordance with an embodiment of the present invention;

Figure 6 is a schematic diagram of an illustrative system with a plurality of components in accordance with an embodiment of the present invention; and

Figure 7 is a schematic diagram of a representative hardware environment in  
5 accordance with an embodiment of the present invention.

### DETAILED DESCRIPTION

Figure 1 is a flowchart of a process 100 for auditing an extensible markup  
10 language (XML) message in a message stream in accordance with an embodiment of the present invention. Messages in a network-based message stream are captured in operation 102 including at least one message in an XML format. Each message in the XML format is then extracted from the captured messages in operation 104 and has a timestamp applied thereto in operation 106. Each  
15 timestamped message in the XML format is then stored in a memory in operation 108.

In one aspect of the present invention, the message stream may include a plurality of messages using a variety of packet-based communication protocols, packaging  
20 standards, transports and formats such as, for example, XML, Transmission Control Protocol (TCP/IP), Hypertext Transfer Protocol (HTTP), File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Serial Line Internet Protocol (SLIP), User Datagram Protocol (UDP), Internetwork Packet Exchange (IPX) Simple Object Access Protocol (SOAP), Multi-Purpose Internet Mail  
25 Extensions (MIME), Java Message Service (JMS). In another aspect, the timestamp may include a digital signature. In a further aspect, the memory may comprise an optical storage medium such as a CD-ROM or DVD-ROM. In an additional aspect, the timestamped message in the XML format may be encrypted prior to storage in the memory.

30

In one embodiment of the present invention, the captured messages may be parsed to identify each message in the XML format for extraction. In another

embodiment, the message stream may be carried out over a communication path (i.e., network) having one or more network segments and where messages are captured at each network segment. In such an embodiment, the captured messages may then be transmitted from each network segment to an aggregation  
5 module prior to extraction of the messages in the XML format.

In a further embodiment of the present invention, a report (e.g., a daily digest) relating to the captured messages may be periodically generated and then transmitted to a remote location to help prevent rollback attacks. In an additional  
10 embodiment, the message stream may transverse a security boundary such as, for example, a firewall having first and second sides where messages on the first side of the security boundary are in an encrypted format and messages on the second side of the security boundary are in an encrypted format. In this embodiment, an encrypted version of each message in the XML format may be captured, extracted,  
15 timestamped, and stored on the first side of the security boundary while a plaintext version of each message in the XML format is captured, extracted, timestamped, and stored on the second side of the security boundary. As a further option, the encrypted and plaintext version of each message in the XML format may then be correlated to detect any changes between the versions of the respective message.

20

The process set forth in Figure 1 requires an understanding of the complete process of interception and analysis including:

1. Capturing - The artifact captures all of the message traffic on the wire.
2. Decoding - The artifact extracts the XML message traffic from the message  
25 stream and determines to the greatest extent possible the transaction to which it corresponds.
3. Timestamping - The artifact must apply a tamperproof timestamp to the message.
4. Archiving - The artifact stores the message to non-volatile media for future  
30 access.

Because the data is generated and processed at a large number of different machines, it may be impractical to modify each program in use. One reasonable procedure is to have a network device that captures the data as it traverses the network. Note that this may require some changes in the network topology to deal  
5 with switches.

In a preferred embodiment, the system incorporates an XML message decoder capable of recognizing various kinds of messages. The requirement here is to reassemble the TCP stream and then determine what kind of message is being  
10 transmitted. The message is then parsed at least enough to determine message identifiers - if any. This may also necessitate pluggable protocol parsing modules.

Once the messages have been captured, they are timestamped and stored. The timestamps should be tamperproof. One approach would simply be to have a very  
15 large hard drive which is tamperproof. Another approach is to use a digital signature on the data and then store the signed messages on an insecure medium.

It should be noted that the device may create a potentially adversarial problem/relationship with the customer in the situation where the customer wishes  
20 to change a message and then reinsert it into the audit trail. A number of technical countermeasures are available to solve this problem. However, it is recommended that the device contain a trusted time source and protect its private key such as, for example, by containing it in a physically tamperproof module. It should also be noted that a single audit device may not be able prevent an attacker from  
25 presenting bogus data to the device during the capture phase. For instance, an attacker might place the capture device on a separate network segment and feed it modified versions of each message. One suitable countermeasure is to have two audit devices, one that operates on a secure network segment and one that operates on an unsecure network segment, then correlate the message traffic sent over the  
30 two segments.

Figure 2 is a schematic diagram of a system 200 for carrying out a process for auditing an XML message in a network-based message stream 202 in accordance with an embodiment of the present invention. This system comprises four separate components including a capture module 204, a parsing module 206, a signature module 208, and a storage module 210 (also referred to as an archive module).

In closer detail, the capture module 204 serves to capture messages in a message stream 202 traversing a network. In a preferred embodiment, the capture module may run on a commercial CPU running a general purpose operating system such as NetBSD. Capture may be possible in at least two ways. For low-load situations, the capture module may use a Berkeley Packet Filter (BPF) to capture all Ethernet traffic and reassemble it in a user space. For high load situations, an enhanced operating system kernel may be utilized. The modification to the kernel allows a socket option that suppresses the output functions. As far as the application is concerned, the kernel accepts connections on a given socket but, preferably, it never transmits. The application then issues read, but not write, commands on the socket to read the application data. In one embodiment, the capture module may include a TCP/IP reassembler to reassemble captured TCP packets into application level message.

Figure 3 is a schematic representation of system 300 for auditing XML messages in a switched network 302 in accordance with an embodiment of the present invention. In a switched network, it may be necessary to have multiple capture devices 304, 306, 308, with one capture device per network segment. The capture devices 304, 306, 308, then transmit their results back to a central module/device 310 which includes for at least one of the other three remaining modules (e.g., at least one of the parsing module, the signature module, and storage module).

Once the module captures messages, the captured messages must be parsed by the parsing module 206 to find their inner XML message for extraction (by the parsing module as well). This may easily be done on the same machine on which



the capture module resides, but in the multiple capture device situation (see Figure 3), one central aggregation and parsing point may be preferred and therefore, the parsing module reside on the central module/device 308. Like the capture module 204, the parsing module 206 may be run on commodity hardware.

5

Figure 4 is a schematic representation of components of a parsing module 206 in accordance with an exemplary embodiment of the present invention. In order to extract XML messages from the captured messages, the parsing device needs to first figure out if the captured messages contain XML messages. Because these  
10 messages may flow over many application level transports, use many different packaging standards, and use many different XML message protocols, etc., getting at this data requires a stackable message unraveler 402. For each potential combination of layers with an XML message at the top, the capture device has a registered stack of pluggable unravelers 404. As illustrated in Figure 4, some  
15 examples of pluggable unravelers 404 that may be plugged into the stack include: an Open Applications Group Integration Specification (OAGIS) pluggable unraveler, a BizTalk pluggable unraveler, a SOAP pluggable unraveler, a MIME pluggable unraveler, a SMTP pluggable unraveler, a TTXML pluggable unraveler, a GCI pluggable unraveler, an Electronic Business XML (ebXML) pluggable  
20 unraveler, a JMS pluggable unraveler, a RosettaNet pluggable unraveler, a MQSeries pluggable unraveler, as well as customizable pluggable unravelers. Also, as option, the parsing module - instead of the capture module - may include the TCP/IP reassembler 406 to reassemble TCP packets captured by the capture module into application level message prior to parsing and extraction.

25

In operation, as messages come in from the TCP reassembler, the unraveling framework 402 and 404 examines the application header to see if it supports that header. It then examines the next level of header and proceeds in this manner until it either finds an unrecognized header or finds the top-level XML message,  
30 which it passes on to a data extractor 408 which extracts the detected XML message.

The signature module 208 is utilized to applied a timestamp to each XML message extracted by the parsing module 206. The signature module preferably connects directly to the parsing module. In a preferred embodiment, the signature module does not run on commodity hardware and instead runs on a special a  
5   tamperproof device containing: a secure time source and secure key storage.

There are at least three possible levels of secure timestamping. A low-level solution has a master clock in a separate network device and individual clocks in each capture device. In one embodiment, the master clock may be a radio clock in  
10   a tamperproof network device that gets its time from the ordinary government time frequencies. The individual clocks may be high quality quartz clocks in the capture devices. At a configurable interval, the master clock establishes a secure channel to the capture devices and interrogates their individual clocks. If an individual clock has drifted beyond a configurable drift window, the master clock  
15   resets it. If the individual clock has a time that is in a configurable danger window, the master clock assumes an attack is under way and sounds an alarm to the management console. In a medium-level solution, each capture device may have its own radio clock in a tamperproof housing. In a really high-level device, each capture device may have its own atomic clock. In all cases, the secure time  
20   source preferably connects to the secure key storage and signature module in a tamperproof fashion.

Essentially, the signature module takes in messages and outputs timestamps. A timestamp is a signed token. In a preferred embodiment, the signature covers: the  
25   digest of the message, the time, a counter which increases by one each signature, and an encrypted digest of all previous messages.

The archive module 210 stores each message to long term storage along with the associated timestamp. Preferably, the long term storage may be write only  
30   medium such as, for example, a large hard drive backed by a DVD writer. Messages may be stored both sequentially and indexed by as much indexing information as the parsing module can extract from the message. It should be

noted that the security of the system does not rely on the write-only properties of the DVD though that may provide additional backup. DVD is merely a cheap high capacity archival storage medium. The archive module 210 may be run on commodity hardware. In one embodiment, the archive module 210 may runs on  
5 the same machine as the parsing module 208.

As a further security measure, the archive module can also send a daily digest of all captured messages (provided by the capture module(s)) to a remote location. This helps to prevent any rollback attack.

10 Optionally, the archive module can encrypt the data before committing it to long term storage. Since public key technology is available, we can encrypt the data under a key not available to any of the online units, and stored in a data key available only to a limited number of authorized users. This allows turnkey  
15 auditability without any long term exposure of the data.

It's becoming increasingly common for electronic commerce operations to be performed over Secured Socket Layer (SSL). In such a case there may be four possibilities:

- 20
1. Encryption performed on server, private key available;
  2. Encryption performed on server, private key unavailable;
  3. Encryption performed on accelerator, private key available; and
  4. Encryption performed on accelerator, private key unavailable.

25 In cases 1 and 3, one can record both the plaintext and the ciphertext and prove that they match. In case 2, one can record the ciphertext only and then demonstrate the plaintext at some later time if provided with the private key. In case 4, one can record both the plaintext and the ciphertext but cannot prove that they match without the private key. As a special case, if the client cooperates with  
30 the audit machine, it can supply the SSL PreMaster Secret, thus enabling decryption even in case 4.

Figure 5 is a schematic representation of an embodiment of a system 500 capable of carrying out multiple correlated capture in accordance with an embodiment of the present invention. As a countermeasure to attacks where the customer forges traffic to the device 200 there is a second recording device 502 (similar to device 5 200) outside the firewall 504 or even on the ISP's side of the network, e.g., between the customer router and the customer's line. So there are two sets of recorded data: one of the plaintext messages inside the enterprise network boundary (see plaintext message stream 506) and one of the SSL encrypted messages outside the enterprise network boundary (see SSL encrypted message stream 508). These can be correlated when it becomes necessary to prove a 10 given piece of a transaction.

Typically, enterprises use an SSL accelerator to decrypt SSL traffic once it crosses the enterprise network boundary. Most of these accelerators preserve the original 15 source IP and port when they decrypt the traffic. Therefore, in order to match up SSL and plaintext data streams a piece of analysis software simply matches up the socket connections. This software can make these matches without having the keying material. Note that this correlation isn't proof of identity, though rough traffic analysis based on record length can provide a certain level of assurance. 20 However, if there is a serious dispute, the parties can reveal their private keys and with the software can take the correlated streams and then decrypt the SSL stream to demonstrate that they are the same.

In one embodiment, the modules (other than the signature module) may be generic 25 motherboards in 19" rack mount chasses. As previously noted, the storage module may need a large drive and some sort of permanent storage device such as a DVD writer or a tape drive. The signature module is preferably specially constructed to be tamperproof. The canonical piece of such technology is the BBN SafeKeyper. In such an embodiment, the signature module should connect 30 directly to a dedicated port on the parsing module. All other modules may simply have standard Ethernet ports.

Figure 6 illustrates an exemplary system 600 with a plurality of components 602 in accordance with one embodiment of the present invention. As shown, such components include a network 604 which take any form including, but not limited to a local area network, a wide area network such as the Internet, and a wireless  
5 network 605. Coupled to the network 604 is a plurality of computers which may take the form of desktop computers 606, lap-top computers 608, hand-held computers 610 (including wireless devices 612 such as wireless PDA's or mobile phones), or any other type of computing hardware/software. As an option, the various computers may be connected to the network 604 by way of a server 614  
10 which may be equipped with a firewall for security purposes. It should be noted that any other type of hardware or software may be included in the system and be considered a component thereof.

A representative hardware environment associated with the various components of  
15 Figure 6 is depicted in Figure 7. In the present description, the various sub-components of each of the components may also be considered components of the system. For example, particular software modules executed on any component of the system may also be considered components of the system. Figure 7 illustrates an illustrative hardware configuration of a workstation 700 having a central  
20 processing unit 702, such as a microprocessor, and a number of other units interconnected via a system bus 704.

The workstation shown in Figure 7 includes a Random Access Memory (RAM) 706, Read Only Memory (ROM) 708, an I/O adapter 710 for connecting  
25 peripheral devices such as, for example, disk storage units 712 and printers 714 to the bus 704, a user interface adapter 716 for connecting various user interface devices such as, for example, a keyboard 718, a mouse 720, a speaker 722, a microphone 724, and/or other user interface devices such as a touch screen or a digital camera to the bus 704, a communication adapter 726 for connecting the  
30 workstation 700 to a communication network 728 (e.g., a data processing network) and a display adapter 730 for connecting the bus 704 to a display device 732.

Also, an article of manufacture, such as a pre-recorded disk or other similar computer program product, for use with a data processing system, could include a storage medium and program means recorded thereon for directing the data processing system to facilitate the practice of the invention. Such apparatus and  
5 articles of manufacture also fall within the spirit and scope of the invention.

A packet is the unit of data that is routed between an origin and a destination on the Internet or any other packet-switch network. When any file (e-mail message, HTML file, Graphics Interchange Format (GIF) file, Uniform Resource Locator  
10 (URL) request, and so forth) is sent from one place to another on the Internet, the Transmission Control Protocol (TCP) layer of TCP/IP divides the file into "chunks" of an efficient size for routing. Each of these packets is separately numbered and includes the Internet address of the destination. The individual packets for a given file may travel different routes through the Internet. When  
15 they have all arrived, they are reassembled into the original file (by the TCP layer at the receiving end).

Packet-switched describes the type of network in which relatively small units of data called packet are routed through a network based on the destination address  
20 contained within each packet. Breaking communication down into packets allows the same data path to be shared among many users in the network. This type of communication between sender and receiver is known as connectionless (rather than dedicated). Most traffic over the Internet uses packet switching and the Internet is basically a connectionless network.

25 Contrasted with packet-switched is circuit-switched, a type of network such as the regular voice telephone network in which the communication circuit (path) for the call is set up and dedicated to the participants in that call. For the duration of the connection, all resources on that circuit are unavailable for other users. Voice  
30 calls using the Internet's packet-switched system are possible. Each end of the conversation is broken down into packets that are reassembled at the other end.

Another common type of digital network that uses packet-switching is the X.25 network, a widely installed commercial wide area network protocol. Internet protocol packets can be carried on an X.25 network. The X.25 network can also support virtual circuits in which a logical connection is established for two parties  
5 on a dedicated basis for some duration. A permanent virtual circuit reserves the path on an ongoing basis and is an alternative for corporations to a system of leased line. A permanent virtual circuit is a dedicated logical connection but the actual physical resources can be shared among multiple logical connections or users.

10

A firewall is a set of related programs, located at a network gateway server, that protects the resources of a private network from users from other networks. (The term also implies the security policy that is used with the programs.) An enterprise with an intranet that allows its workers access to the wider Internet  
15 installs a firewall to prevent outsiders from accessing its own private data resources and for controlling what outside resources its own users have access to.

Basically, a firewall, working closely with a router program, examines each network packet to determine whether to forward it toward its destination. A  
20 firewall also includes or works with a proxy server that makes network requests on behalf of workstation users. A firewall is often installed in a specially designated computer separate from the rest of the network so that no incoming request can get directly at private network resources.

25 There are a number of firewall screening methods. A simple one is to screen requests to make sure they come from acceptable (previously identified) domain name and Internet Protocol (IP) addresses. For mobile users, firewalls allow remote access in to the private network by the use of secure logon procedures and authentication certificates.

30

Common features of firewall include logging and reporting, automatic alarms at given thresholds of attack, and a graphical user interface for controlling the firewall.

- 5 Transmission Control Protocol/Internet Protocol (TCP/IP) is a basic communication language or protocol of the Internet. It can also be used as a communications protocol in the private networks called intranet and in extranet. When you are set up with direct access to the Internet, your computer is provided with a copy of the TCP/IP program just as every other computer that you may  
10 send messages to or get information from also has a copy of TCP/IP.

- TCP/IP is a two-layering program. The higher layer, Transmission Control Protocol (TCP), manages the assembling of a message or file into smaller packet that are transmitted over the Internet and received by a TCP layer that reassembles  
15 the packets into the original message. The lower layer, Internet Protocol (IP), handles the address part of each packet so that it gets to the right destination. Each gateway computer on the network checks this address to see where to forward the message. Even though some packets from the same message are routed differently than others, they'll be reassembled at the destination.

- 20 TCP/IP uses a client/server model of communication in which a computer user (a client) requests and is provided a service (such as sending a Web page) by another computer (a server) in the network. TCP/IP communication is primarily point-to-point, meaning each communication is from one point (or host computer) in the  
25 network to another point or host computer. TCP/IP and the higher-level applications that use it are collectively said to be "stateless" because each client request is considered a new request unrelated to any previous one (unlike ordinary phone conversations that require a dedicated connection for the call duration). Being stateless frees network paths so that everyone can use them continuously.  
30 (Note that the TCP layer itself is not stateless as far as any one message is concerned. Its connection remains in place until all packets in a message have been received.).



Many Internet users are familiar with the even higher layer application protocols that use TCP/IP to get to the Internet. These include the World Wide Web's Hypertext Transfer Protocol (HTTP), the File Transfer Protocol (FTP), Telnet  
5 which lets you logon to remote computers, and the Simple Mail Transfer Protocol (SMTP). These and other protocols are often packaged together with TCP/IP as a "suite."

Personal computer users usually get to the Internet through the Serial Line Internet Protocol (SLIP) or the Point-to-Point Protocol. These protocols encapsulate the IP  
10 packets so that they can be sent over a dial-up phone connection to an access provider's modem.

Protocols related to TCP/IP include the User Datagram Protocol (UDP), which is used instead of TCP for special purposes. Other protocols are used by network  
15 host computers for exchanging router information. These include the Internet Control Message Protocol (ICMP), the Interior Gateway Protocol (IGP), the Exterior Gateway Protocol (EGP), and the Border Gateway Protocol (BGP).

Internetwork Packet Exchange (IPX) is a networking protocol from Novell that  
20 interconnects networks that use Novell's NetWare clients and servers. IPX is a datagram or packet protocol. IPX works at the network layer of communication protocols and is connectionless (that is, it doesn't require that a connection be maintained during an exchange of packets as, for example, a regular voice phone call does).

25 Packet acknowledgment is managed by another Novell protocol, the Sequenced Packet Exchange (SPX). Other related Novell NetWare protocols are: the Routing Information Protocol (RIP), the Service Advertising Protocol (SAP), and the NetWare Link Services Protocol (NLSP).

30 A virtual private network (VPN) is a private data network that makes use of the public telecommunication infrastructure, maintaining privacy through the use of a

tunneling protocol and security procedures. A virtual private network can be contrasted with a system of owned or leased lines that can only be used by one company. The idea of the VPN is to give the company the same capabilities at much lower cost by using the shared public infrastructure rather than a private one. Phone companies have provided secure shared resources for voice messages. A virtual private network makes it possible to have the same secure sharing of public resources for data.

Using a virtual private network involves encryption data before sending it through the public network and decrypting it at the receiving end. An additional level of security involves encrypting not only the data but also the originating and receiving network addresses. Microsoft, 3Com, and several other companies have developed the Point-to-Point Tunneling Protocol (PPP) and Microsoft has extended Windows NT [http://www.whatis.com/WhatIs\\_Definition\\_Page/0,4152,213368,00.html](http://www.whatis.com/WhatIs_Definition_Page/0,4152,213368,00.html) to support it. VPN software is typically installed as part of a company's firewall server.

XML (Extensible Markup Language) is a flexible way to create common information formats and share both the format and the data on the World Wide Web, intranets, and elsewhere. For example, computer makers might agree on a standard or common way to describe the information about a computer product (processor speed, memory size, and so forth) and then describe the product information format with XML. Such a standard way of describing data would enable a user to send an intelligent agent (a program) to each computer maker's Web site, gather data, and then make a valid comparison. XML can be used by any individual or group of individuals or companies that wants to share information in a consistent way.

XML, a formal recommendation from the World Wide Web Consortium (W3C), is similar to the language of today's Web pages, the Hypertext Markup Language (HTML). Both XML and HTML contain markup symbols to describe the

contents of a page or file. HTML, however, describes the content of a Web page (mainly text and graphic images) only in terms of how it is to be displayed and interacted with. For example, the letter "p" placed within markup tags starts a new paragraph. XML describes the content in terms of what data is being

5 described. For example, the word "phonenum" placed within markup tags could indicate that the data that followed was a phone number. This means that an XML file can be processed purely as data by a program or it can be stored with similar data on another computer or, like an HTML file, that it can be displayed. For example, depending on how the application in the receiving computer wanted to

10 handle the phone number, it could be stored, displayed, or dialed.

XML is "extensible" because, unlike HTML, the markup symbols are unlimited and self-defining. XML is actually a simpler and easier-to-use subset of the Standard Generalized Markup Language (SGML), the standard for how to create a

15 document structure. It is expected that HTML and XML will be used together in many Web applications. XML markup, for example, may appear within an HTML page.

Early applications of XML include Microsoft's Channel Definition Format (CDF),

20 which describes a channel, a portion of a Web site that has been downloaded to your hard disk and is then is updated periodically as information changes. A specific CDF file contains data that specifies an initial Web page and how frequently it is updated. Another early application is ChartWare, which uses XML as a way to describe medical charts so that they can be shared by doctors.

25 Applications related to banking, e-commerce ordering, personal preference profiles, purchase orders, litigation documents, part lists, and many others are anticipated.

On the Internet, B2B (business-to-business), also known as e-biz, is the exchange

30 of products, services, or information between businesses rather than between businesses and consumers.

Encryption is the conversion of data into a form, called a ciphertext, that cannot be easily understood by unauthorized people. Decryption is the process of converting encrypted data back into its original form, so it can be understood.

5 The use of encryption/decryption is as old as the art of communication. In wartime, a cipher, often incorrectly called a "code," can be employed to keep the enemy from obtaining the contents of transmissions (technically, a code is a means of representing a signal without the intent of keeping it secret; examples are Morse code and

10 ASCII [http://www.whatis.com/WhatIs\\_Definition\\_Page/0,4152,211600,00.html](http://www.whatis.com/WhatIs_Definition_Page/0,4152,211600,00.html)). Simple ciphers include the substitution of letters for numbers, the rotation of letters in the alphabet, and the "scrambling" of voice signals by inverting the sideband frequencies. More complex ciphers work according to sophisticated computer algorithm that rearrange the data bits in digital signals.

15 In order to easily recover the contents of an encrypted signal, the correct decryption key is required. The key is an algorithm that "undoes" the work of the encryption algorithm. Alternatively, a computer can be used in an attempt to "break" the cipher. The more complex the encryption algorithm, the more  
20 difficult it becomes to eavesdrop on the communications without access to the key.

Rivest-Shamir-Adleman (RSA) is an Internet encryption and authentication system that uses an algorithm developed in 1977 by Ron Rivest, Adi Shamir, and  
25 Leonard Adleman. The RSA algorithm is a commonly used encryption and authentication algorithm and is included as part of the Web browser from Netscape and Microsoft. It's also part of Lotus Notes, Intuit's Quicken, and many other products. The encryption system is owned by RSA Security.

30 The RSA algorithm involves multiplying two large prime numbers (a prime number is a number divisible only by that number and 1) and through additional operations deriving a set of two numbers that constitutes the public key and

another set that is the private key. Once the keys have been developed, the original prime numbers are no longer important and can be discarded. Both the public and the private keys are needed for encryption /decryption but only the owner of a private key ever needs to know it. Using the RSA system, the private  
5 key never needs to be sent across the Internet.

The private key is used to decrypt text that has been encrypted with the public key. Thus, if I send you a message, I can find out your public key (but not your private key) from a central administrator and encrypt a message to you using your  
10 public key. When you receive it, you decrypt it with your private key. In addition to encrypting messages (which ensures privacy), you can authenticate yourself to me (so I know that it is really you who sent the message) by using your private key to encrypt a digital certificate. When I receive it, I can use your public key to decrypt it.

15 Secure Sockets Layer (SSL) is a commonly-used protocol for managing the security of a message transmission on the Internet. SSL uses a program layer located between the Internet's Hypertext Transfer Protocol (HTTP) and Transport Control Protocol (TCP) layers. The "sockets" part of the term refers to the sockets  
20 method of passing data back and forth between a client and a server program in a network or between program layers in the same computer. SSL uses the public-and-private key encryption system from RSA, which also includes the use of a digital certificate.

25 A digital signature is an electronic rather than a written signature that can be used by someone to authenticate the identity of the sender of a message or of the signer of a document. It can also be used to ensure that the original content of the message or document that has been conveyed is unchanged. Additional benefits to the use of a digital signature are that it is easily transportable, cannot be easily  
30 repudiated, cannot be imitated by someone else, and can be automatically time-stamped.

A digital signature can be used with any kind of message, whether it is encrypted or not, simply so that the receiver can be sure of the sender's identity and that the message arrived intact. A digital certificate contains the digital signature of the certificate-issuing authority so that anyone can verify that the certificate is real.

5

BizTalk is an industry initiative headed by Microsoft to promote Extensible Markup Language (XML) as the common data exchange language for e-commerce and application integration on the Internet. While not a standards body  
10 per se, the group is fostering a common XML message-passing architecture to tie systems together. BizTalk says that the growth of e-commerce requires businesses using different computer technologies to have a means to share data. Accepting XML as a platform-neutral way to represent data transmitted between computers, the BizTalk group provides guidelines, referred to as the BizTalk  
15 Framework, for how to publish schema (standard data structures) in XML and how to use XML messages to integrate software programs.

Simple Object Access Protocol (SOAP) is a way for a program running in one kind of operating system to communicate with a program in the same or another  
20 kind of an operating system by using the World Wide Web's Hypertext Transfer Protocol and its Extensible Markup Language (XML) as the mechanisms for information exchange. Since Web protocol are installed and available for use by all major operating system platforms, HTTP and XML provide an already at-hand solution to the problem of how programs running under different operating  
25 systems in a network can communicate with each other. SOAP specifies exactly how to encode an HTTP header and an XML file so that a program in one computer can call a program in another computer and pass it information. It also specifies how the called program can return a response.

30 SOAP was developed by Microsoft, DevelopMentor, and Userland Software and has been proposed as a standard interface to the Internet Engineering Task Force (IETF). It is somewhat similar to the Internet Inter-ORB Protocol, a protocol that

is part of the Common Object Request Broker Architecture. Sun Microsystems' Remote Method Invocation is a similar client/server <http://whatIs.techtarget.com/WhatIs/Search/Results/Exact/clientse.htm> interprogram protocol between programs written in Java.

5

An advantage of SOAP is that program calls are much more likely to get through firewall servers that screen out requests other than those for known applications (through the designated port mechanism). Since HTTP requests are usually allowed through firewalls, programs using SOAP to communicate can be sure that they can communicate with programs anywhere.

10

Multi-Purpose Internet Mail Extensions (MIME) is an extension of the original Internet e-mail protocol that lets people use the protocol to exchange different kinds of data files on the Internet: audio, video, images, application programs, and other kinds, as well as the ASCII handled in the original protocol, the Simple Mail Transport Protocol (SMTP). In 1991, Nathan Borenstein of Bellcore proposed to the IETF that SMTP be extended so that Internet (but mainly Web) client and server could recognize and handle other kinds of data than ASCII text. As a result, new file types were added to "mail" as a supported Internet Protocol file type.

15  
20

Servers insert the MIME header at the beginning of any Web transmission. Clients use this header to select an appropriate "player" application for the type of data the header indicates. Some of these players are built into the Web client or browser (for example, all browser come with GIF and JPEG image players as well as the ability to handle HTML files); other players may need to be downloaded.

25

New MIME data types are registered with the Internet Assigned Numbers Authority (IANA).

30

MIME is specified in detail in Internet Request for Comments 1521 and 1522, which amend the original mail protocol specification, RFC 821 (the Simple Mail Transport Protocol) and the ASCII messaging header, RFC 822.

- 5 MQSeries is an IBM software family whose components are used to tie together other software applications so that they can work together. This type of application is often known as business integration software or middleware.

MQSeries consists of three products:

- 10
- MQSeries Messaging, which provides the communication mechanism between applications on different platforms
  - MQSeries Integrator, which centralizes and applies business operations rules
  - MQSeries Workflow, which enables the capture, visualization, and
- 15 automation of business processes

- The point of business integration is to connect different computer systems, diverse geographical locations, and dissimilar IT infrastructures so that a seamless operation can be run. IBM's MQSeries supplies communications between applications, or between users and a set of applications on dissimilar systems. It
- 20 has grown in popularity as applications are made available over the Internet because of its support of over 35 platforms and its ability to integrate disparate automation systems.

- An additional helpful feature is that its messaging scheme requires the application
- 25 that receives the message to confirm receipt. If no confirmation materializes, the message is re-sent by the MQSeries.

- Java Message Service (JMS) is an application program interface from Sun Microsystems that supports the formal communication known as messaging
- 30 between computers in a network. Sun's JMS provides a common interface to standard messaging protocols and also to special messaging services in support of Java programs.



The messages involved exchange crucial data between computers - rather than between users - and contain information such as event notification and service requests. Messaging is often used to coordinate programs in dissimilar systems or  
5 written in different programming languages.

Using the JMS interface, a programmer can invoke the messaging services of IBM's MQSeries, Progress Software's SonicMQ, and other popular messaging product vendors. In addition, JMS supports messages that contain serialized Java  
10 object and messages that contain Extensible Markup Language (XML) pages.

RosettaNet is an organization set up by leading information technology companies to define and implement a common set of standards for e-business. RosettaNet is defining a common parts dictionary so that different companies can define the  
15 same product the same way. It is also defining up to 100 e-business transaction processes and standardizing them. Because RosettaNet is supported by all or most of the major companies in the IT industry, its standards are expected to be widely adopted.

20 RosettaNet has developed a structured four-part approach for creating what it calls Partner Interface Processes (PIPs).

- Business Process Modeling examines common business procedures and defines the components of the processes.
- Business Process Analysis analyzes the processes and defines a target list  
25 of desirable changes to the processes.
- PIP Development establishes guidelines and documentation for the changes.
- Dictionaries consist of two data dictionary: a technical properties dictionary and a business properties dictionary. Along with the RosettaNet  
30 Implementation Framework (which defines an exchange protocol for PIP implementation), the dictionaries form the basis for PIP development.

- Simple Mail Transfer Protocol (SMTP) is a TCP/IP protocol used in sending and receiving e-mail. However, since it's limited in its ability to queue messages at the receiving end, it's usually used with one of two other protocols, POP3 or Internet Message Access Protocol, that let the user save messages in a server mailbox and download them periodically from the server. In other words, users typically use a program that uses SMTP for sending e-mail and either POP3 or IMAP for receiving messages that have been received for them at their local server. Most mail programs such as Eudora let you specify both an SMTP server and a POP server. On
- 10 UNIX [http://whatis.techtarget.com/WhatIs\\_Definition\\_Page/0,4152,213253,00.html](http://whatis.techtarget.com/WhatIs_Definition_Page/0,4152,213253,00.html)-based systems, sendmail is the most widely-used SMTP server for e-mail. A commercial package, Sendmail, includes a POP3 server and also comes in a version for Windows NT.
- 15 SMTP usually is implemented to operate over Transmission Control Protocol port 25. The details of SMTP are in Request for Comments 821 of the Internet Engineering Task Force (IETF). An alternative to SMTP that is widely used in Europe is X400.
- 20 The Hypertext Transfer Protocol (HTTP) is the set of rules for exchanging files (text, graphic images, sound, video, and other multimedia files) on the World Wide Web. Relative to the TCP/IP suite of protocols (which are the basis for information exchange on the Internet), HTTP is an application protocol.
- 25 Essential concepts that are part of HTTP include (as its name implies) the idea that files can contain references to other files whose selection will elicit additional transfer requests. Any Web
- server [http://whatis.techtarget.com/WhatIs\\_Definition\\_Page/0,4152,213606,00.html](http://whatis.techtarget.com/WhatIs_Definition_Page/0,4152,213606,00.html) machine contains, in addition to the HTML and other files it can serve, an
- 30 HTTP daemon, a program that is designed to wait for HTTP requests and handle them when they arrive. Your Web browser is an HTTP
- client [http://whatis.techtarget.com/WhatIs\\_Definition\\_Page/0,4152,211795,00.html](http://whatis.techtarget.com/WhatIs_Definition_Page/0,4152,211795,00.html)

1 sending requests to server machines. When the browser user enters file requests by either "opening" a Web file (typing in a URL) or clicking on a hypertext link, the browser builds an HTTP request and sends it to the Internet Protocol address indicated by the URL. The HTTP daemon in the destination server machine  
5 receives the request and, after any necessary processing, the requested file is returned.

An embodiment of the present invention may also be written using Java, C, and the C++ language and utilize object oriented programming methodology. Object  
10 oriented programming (OOP) has become increasingly used to develop complex applications. As OOP moves toward the mainstream of software design and development, various software solutions require adaptation to make use of the benefits of OOP. A need exists for these principles of OOP to be applied to a messaging interface of an electronic messaging system such that a set of OOP  
15 classes and objects for the messaging interface can be provided.

OOP is a process of developing computer software using objects, including the steps of analyzing the problem, designing the system, and constructing the program. An object is a software package that contains both data and a collection  
20 of related structures and procedures. Since it contains both data and a collection of structures and procedures, it can be visualized as a self-sufficient component that does not require other additional structures, procedures or data to perform its specific task. OOP, therefore, views a computer program as a collection of largely autonomous components, called objects, each of which is responsible for a  
25 specific task. This concept of packaging data, structures, and procedures together in one component or module is called encapsulation.

In general, OOP components are reusable software modules which present an interface that conforms to an object model and which are accessed at run-time  
30 through a component integration architecture. A component integration architecture is a set of architecture mechanisms which allow software modules in different process spaces to utilize each others capabilities or functions. This is

generally done by assuming a common component object model on which to build the architecture. It is worthwhile to differentiate between an object and a class of objects at this point. An object is a single instance of the class of objects, which is often just called a class. A class of objects can be viewed as a blueprint, from  
5 which many objects can be formed.

OOP allows the programmer to create an object that is a part of another object. For example, the object representing a piston engine is said to have a composition-relationship with the object representing a piston. In reality, a piston engine  
10 comprises a piston, valves and many other components; the fact that a piston is an element of a piston engine can be logically and semantically represented in OOP by two objects.

OOP also allows creation of an object that “depends from” another object. If  
15 there are two objects, one representing a piston engine and the other representing a piston engine wherein the piston is made of ceramic, then the relationship between the two objects is not that of composition. A ceramic piston engine does not make up a piston engine. Rather it is merely one kind of piston engine that has one more limitation than the piston engine; its piston is made of ceramic. In this case,  
20 the object representing the ceramic piston engine is called a derived object, and it inherits all of the aspects of the object representing the piston engine and adds further limitation or detail to it. The object representing the ceramic piston engine “depends from” the object representing the piston engine. The relationship between these objects is called inheritance.

25 When the object or class representing the ceramic piston engine inherits all of the aspects of the objects representing the piston engine, it inherits the thermal characteristics of a standard piston defined in the piston engine class. However, the ceramic piston engine object overrides these ceramic specific thermal  
30 characteristics, which are typically different from those associated with a metal piston. It skips over the original and uses new functions related to ceramic pistons. Different kinds of piston engines have different characteristics, but may

have the same underlying functions associated with it (e.g., how many pistons in the engine, ignition sequences, lubrication, etc.). To access each of these functions in any piston engine object, a programmer would call the same functions with the same names, but each type of piston engine may have

5 different/overriding implementations of functions behind the same name. This ability to hide different implementations of a function behind the same name is called polymorphism and it greatly simplifies communication among objects.

With the concepts of composition-relationship, encapsulation, inheritance and

10 polymorphism, an object can represent just about anything in the real world. In fact, one's logical perception of the reality is the only limit on determining the kinds of things that can become objects in object-oriented software. Some typical categories are as follows:

- Objects can represent physical objects, such as automobiles in a traffic-  
15 flow simulation, electrical components in a circuit-design program, countries in an economics model, or aircraft in an air-traffic-control system.
- Objects can represent elements of the computer-user environment such as windows, menus or graphics objects.
- 20 • An object can represent an inventory, such as a personnel file or a table of the latitudes and longitudes of cities.
- An object can represent user-defined data types such as time, angles, and complex numbers, or points on the plane.

25 With this enormous capability of an object to represent just about any logically separable matters, OOP allows the software developer to design and implement a computer program that is a model of some aspects of reality, whether that reality is a physical entity, a process, a system, or a composition of matter. Since the object can represent anything, the software developer can create an object which

30 can be used as a component in a larger software project in the future.

If 90% of a new OOP software program consists of proven, existing components made from preexisting reusable objects, then only the remaining 10% of the new software project has to be written and tested from scratch. Since 90% already came from an inventory of extensively tested reusable objects, the potential domain from which an error could originate is 10% of the program. As a result, OOP enables software developers to build objects out of other, previously built objects.

This process closely resembles complex machinery being built out of assemblies and sub-assemblies. OOP technology, therefore, makes software engineering more like hardware engineering in that software is built from existing components, which are available to the developer as objects. All this adds up to an improved quality of the software as well as an increased speed of its development.

Programming languages are beginning to fully support the OOP principles, such as encapsulation, inheritance, polymorphism, and composition-relationship. With the advent of the C++ language, many commercial software developers have embraced OOP. C++ is an OOP language that offers a fast, machine-executable code. Furthermore, C++ is suitable for both commercial-application and systems-programming projects. For now, C++ appears to be the most popular choice among many OOP programmers, but there is a host of other OOP languages, such as Smalltalk, Common Lisp Object System (CLOS), and Eiffel. Additionally, OOP capabilities are being added to more traditional popular computer programming languages such as Pascal.

The benefits of object classes can be summarized, as follows:

- Objects and their corresponding classes break down complex programming problems into many smaller, simpler problems.
- Encapsulation enforces data abstraction through the organization of data into small, independent objects that can communicate with each other. Encapsulation protects the data in an object from accidental damage, but

allows other objects to interact with that data by calling the object's member functions and structures.

- Subclassing and inheritance make it possible to extend and modify objects through deriving new kinds of objects from the standard classes available in the system. Thus, new capabilities are created without having to start from scratch.
- Polymorphism and multiple inheritance make it possible for different programmers to mix and match characteristics of many different classes and create specialized objects that can still work with related objects in predictable ways.
- Class hierarchies and containment hierarchies provide a flexible mechanism for modeling real-world objects and the relationships among them.
- Libraries of reusable classes are useful in many situations, but they also have some limitations. For example:
  - Complexity. In a complex system, the class hierarchies for related classes can become extremely confusing, with many dozens or even hundreds of classes.
  - Flow of control. A program written with the aid of class libraries is still responsible for the flow of control (i.e., it must control the interactions among all the objects created from a particular library). The programmer has to decide which functions to call at what times for which kinds of objects.
  - Duplication of effort. Although class libraries allow programmers to use and reuse many small pieces of code, each programmer puts those pieces together in a different way. Two different programmers can use the same set of class libraries to write two programs that do exactly the same thing but whose internal structure (i.e., design) may be quite different, depending on hundreds of small decisions each programmer makes along the way. Inevitably, similar pieces of code end up doing similar things in slightly different ways and do not work as well together as they should.

Class libraries are very flexible. As programs grow more complex, more programmers are forced to reinvent basic solutions to basic problems over and over again. A relatively new extension of the class library concept is to have a framework of class libraries. This framework is more complex and consists of significant collections of collaborating classes that capture both the small scale patterns and major mechanisms that implement the common requirements and design in a specific application domain. They were first developed to free application programmers from the chores involved in displaying menus, windows, dialog boxes, and other standard user interface elements for personal computers.

10

Frameworks also represent a change in the way programmers think about the interaction between the code they write and code written by others. In the early days of procedural programming, the programmer called libraries provided by the operating system to perform certain tasks, but basically the program executed down the page from start to finish, and the programmer was solely responsible for the flow of control. This was appropriate for printing out paychecks, calculating a mathematical table, or solving other problems with a program that executed in just one way.

20 The development of graphical user interfaces began to turn this procedural programming arrangement inside out. These interfaces allow the user, rather than program logic, to drive the program and decide when certain actions should be performed. Today, most personal computer software accomplishes this by means of an event loop which monitors the mouse, keyboard, and other sources of external events and calls the appropriate parts of the programmer's code according to actions that the user performs. The programmer no longer determines the order in which events occur. Instead, a program is divided into separate pieces that are called at unpredictable times and in an unpredictable order. By relinquishing control in this way to users, the developer creates a program that is much easier to use. Nevertheless, individual pieces of the program written by the developer still call libraries provided by the operating system to accomplish certain tasks, and the

30



programmer must still determine the flow of control within each piece after it's called by the event loop. Application code still "sits on top of" the system.

Even event loop programs require programmers to write a lot of code that should  
5 not need to be written separately for every application. The concept of an application framework carries the event loop concept further. Instead of dealing with all the nuts and bolts of constructing basic menus, windows, and dialog boxes and then making these things all work together, programmers using application frameworks start with working application code and basic user  
10 interface elements in place. Subsequently, they build from there by replacing some of the generic capabilities of the framework with the specific capabilities of the intended application.

Application frameworks reduce the total amount of code that a programmer has to  
15 write from scratch. However, because the framework is really a generic application that displays windows, supports copy and paste, and so on, the programmer can also relinquish control to a greater degree than event loop programs permit. The framework code takes care of almost all event handling and flow of control, and the programmer's code is called only when the framework  
20 needs it (e.g., to create or manipulate a proprietary data structure).

A programmer writing a framework program not only relinquishes control to the user (as is also true for event loop programs), but also relinquishes the detailed flow of control within the program to the framework. This approach allows the  
25 creation of more complex systems that work together in interesting ways, as opposed to isolated programs, having custom code, being created over and over again for similar problems.

Thus, as is explained above, a framework basically is a collection of cooperating  
30 classes that make up a reusable design solution for a given problem domain. It typically includes objects that provide default behavior (e.g., for menus and windows), and programmers use it by inheriting some of that default behavior and

overriding other behavior so that the framework calls application code at the appropriate times.

There are three main differences between frameworks and class libraries:

- 5     •     Behavior versus protocol. Class libraries are essentially collections of behaviors that you can call when you want those individual behaviors in your program. A framework, on the other hand, provides not only behavior but also the protocol or set of rules that govern the ways in which behaviors can be combined, including rules for what a programmer is  
10           supposed to provide versus what the framework provides.
- Call versus override. With a class library, the code the programmer instantiates objects and calls their member functions. It's possible to instantiate and call objects in the same way with a framework (i.e., to treat the framework as a class library), but to take full advantage of a  
15           framework's reusable design, a programmer typically writes code that overrides and is called by the framework. The framework manages the flow of control among its objects. Writing a program involves dividing responsibilities among the various pieces of software that are called by the framework rather than specifying how the different pieces should work  
20           together.
- Implementation versus design. With class libraries, programmers reuse only implementations, whereas with frameworks, they reuse design. A framework embodies the way a family of related programs or pieces of software work. It represents a generic design solution that can be adapted  
25           to a variety of specific problems in a given domain. For example, a single framework can embody the way a user interface works, even though two different user interfaces created with the same framework might solve quite different interface problems.
- 30     Thus, through the development of frameworks for solutions to various problems and programming tasks, significant reductions in the design and development effort for software can be achieved. A preferred embodiment of the invention

utilizes HyperText Markup Language (HTML) to implement documents on the Internet together with a general-purpose secure communication protocol for a transport medium between the client and the server. HTTP or other protocols could be readily substituted for HTML without undue experimentation.

- 5 Information on these products is available in T. Berners-Lee, D. Connolly, "RFC 1866: Hypertext Markup Language - 2.0" (Nov. 1995); and R. Fielding, H. Frystyk, T. Berners-Lee, J. Gettys and J.C. Mogul, "Hypertext Transfer Protocol -- HTTP/1.1: HTTP Working Group Internet Draft" (May 2, 1996). HTML is a simple data format used to create hypertext documents that are portable from one
- 10 platform to another. HTML documents are SGML documents with generic semantics that are appropriate for representing information from a wide range of domains. HTML has been in use by the World-Wide Web global information initiative since 1990. HTML is an application of ISO Standard 8879; 1986 Information Processing Text and Office Systems; Standard Generalized Markup
- 15 Language (SGML).

To date, Web development tools have been limited in their ability to create dynamic Web applications which span from client to server and interoperate with existing computing resources. Until recently, HTML has been the dominant

20 technology used in development of Web-based solutions. However, HTML has proven to be inadequate in the following areas:

- Poor performance;
- Restricted user interface capabilities;
- Can only produce static Web pages;
- 25 • Lack of interoperability with existing applications and data; and
- Inability to scale.

Sun Microsystems's Java language solves many of the client-side problems by:

- Improving performance on the client side;
- 30 • Enabling the creation of dynamic, real-time Web applications; and
- Providing the ability to create a wide variety of user interface components.

With Java, developers can create robust User Interface (UI) components. Custom "widgets" (e.g., real-time stock tickers, animated icons, etc.) can be created, and client-side performance is improved. Unlike HTML, Java supports the notion of client-side validation, offloading appropriate processing onto the client for improved performance. Dynamic, real-time Web pages can be created. Using the above-mentioned custom UI components, dynamic Web pages can also be created.

Sun's Java language has emerged as an industry-recognized language for "programming the Internet." Sun defines Java as: "a simple, object-oriented, distributed, interpreted, robust, secure, architecture-neutral, portable, high-performance, multithreaded, dynamic, buzzword-compliant, general-purpose programming language. Java supports programming for the Internet in the form of platform-independent Java applets." Java applets are small, specialized applications that comply with Sun's Java Application Programming Interface (API) allowing developers to add "interactive content" to Web documents (e.g., simple animations, page adornments, basic games, etc.). Applets execute within a Java-compatible browser (e.g., Netscape Navigator) by copying code from the server to client. From a language standpoint, Java's core feature set is based on C++. Sun's Java literature states that Java is basically, "C++ with extensions from Objective C for more dynamic method resolution."

Another technology that provides similar function to Java is provided by Microsoft and ActiveX Technologies, to give developers and Web designers wherewithal to build dynamic content for the Internet and personal computers. ActiveX includes tools for developing animation, 3-D virtual reality, video and other multimedia content. The tools use Internet standards, work on multiple platforms, and are being supported by over 100 companies. The group's building blocks are called ActiveX Controls, small, fast components that enable developers to embed parts of software in hypertext markup language (HTML) pages. ActiveX Controls work with a variety of programming languages including Microsoft Visual C++, Borland Delphi, Microsoft Visual Basic programming

system and, in the future, Microsoft's development tool for Java, code named "Jakarta." ActiveX Technologies also includes ActiveX Server Framework, allowing developers to create server applications. One of ordinary skill in the art readily recognizes that ActiveX could be substituted for Java without undue experimentation to practice the invention.

Based on the foregoing specification, the invention may be implemented using computer programming or engineering techniques including computer software, firmware, hardware or any combination or subset thereof. Any such resulting program, having computer-readable code means, may be embodied or provided within one or more computer-readable media, thereby making a computer program product, i.e., an article of manufacture, according to the invention. The computer readable media may be, for instance, a fixed (hard) drive, diskette, optical disk, magnetic tape, semiconductor memory such as read-only memory (ROM), etc., or any transmitting/receiving medium such as the Internet or other communication network or link. The article of manufacture containing the computer code may be made and/or used by executing the code directly from one medium, by copying the code from one medium to another medium, or by transmitting the code over a network.

One skilled in the art of computer science will easily be able to combine the software created as described with appropriate general purpose or special purpose computer hardware to create a computer system or computer sub-system embodying the method of the invention.

While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

## CLAIMS

What is claimed is:

1. A method for auditing a message in a message stream, comprising:
  - a) capturing messages in a message stream, wherein the messages include at  
5 least one message in an extensible markup language (XML) format;
  - b) extracting the at least one message in the XML format from the captured  
messages;
  - c) applying a timestamp to the extracted at least one message in the XML  
format; and
  - 10 d) storing the timestamped at least one message in the XML format in a  
memory.
2. The method of claim 1, wherein the message stream comprises a plurality  
15 of messages utilizing a plurality of protocols.
3. The method of claim 1, wherein the captured messages are parsed to  
identify the at least one message in the XML format for extraction.
4. The method of claim 1, wherein the message stream is carried out over a  
20 communication path having one or more segments, and wherein messages  
are captured at each segment.
5. The method of claim 4, wherein the captured messages are transmitted  
from each segment to an aggregation module prior to extraction of the at  
25 least one message in the XML format.
6. The method of claim 1, wherein the timestamp includes a digital signature.
7. The method of claim 1, wherein the memory comprises a write once  
30 storage medium.

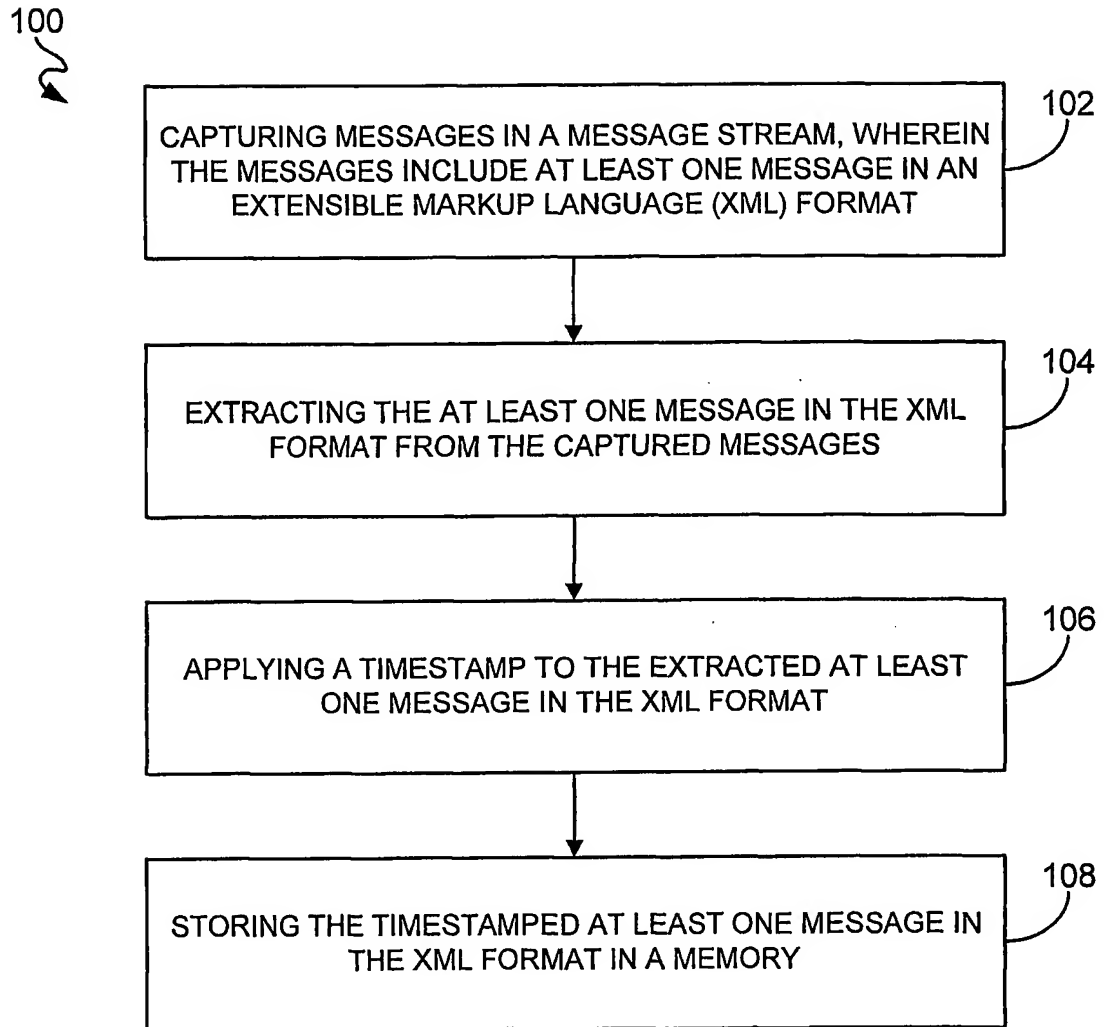
8. The method of claim 1, wherein a report relating to the captured messages is generated.
9. The method of claim 1, wherein the timestamped at least one message in  
5 the XML format is encrypted prior to storage in the memory.
10. The method of claim 1, wherein the message stream transverses a security boundary having first and second sides, wherein messages in the message stream on the first side of the security boundary are in an encrypted  
10 format, wherein messages in the message stream on the second side of the security boundary are in an encrypted format, wherein an encrypted version of the at least one message in the XML format is captured, extracted, and timestamped on the first side of the security boundary, and wherein a plaintext version of the at least one message in the XML format  
15 is captured, extracted, and timestamped on the second side of the security boundary.
11. The method of claim 10, wherein the encrypted and plaintext version of the at least one message in the XML format are correlated to detect any  
20 changes between the versions of the at least one message in the XML format.
12. A system for auditing a message in a message stream, comprising:
  - a) logic for capturing messages in a message stream, wherein the messages  
25 include at least one message in an extensible markup language (XML) format;
  - b) logic for extracting the at least one message in the XML format from the captured messages;
  - c) logic for applying a timestamp to the extracted at least one message in the  
30 XML format; and
  - d) logic for storing the timestamped at least one message in the XML format in a memory.

13. The system of claim 12, wherein the message stream comprises a plurality of message utilizing a plurality of protocols.
- 5 14. The system of claim 12, wherein the captured messages are parsed to identify the at least one message in the XML format for extraction.
15. The system of claim 12, wherein the timestamp includes a digital signature.
- 10 16. The system of claim 12, wherein the timestamped at least one message in the XML format is encrypted prior to storage in the memory.
- 15 17. A computer program product for auditing a message in a message stream, comprising:
- a) computer code for capturing messages in a message stream, wherein the messages include at least one message in an extensible markup language (XML) format;
- 20 b) computer code for extracting the at least one message in the XML format from the captured messages;
- c) computer code for applying a timestamp to the extracted at least one message in the XML format; and
- d) computer code for storing the timestamped at least one message in the XML format in a memory.
- 25 18. The computer program product of claim 17, wherein the message stream comprises a plurality of message utilizing a plurality of protocols.
- 30 19. The computer program product of claim 17, wherein the timestamp includes a digital signature.



20. The computer program product of claim 17, wherein the timestamped at least one message in the XML format is encrypted prior to storage in the memory.

1/7

**FIG. 1**




FIG. 2


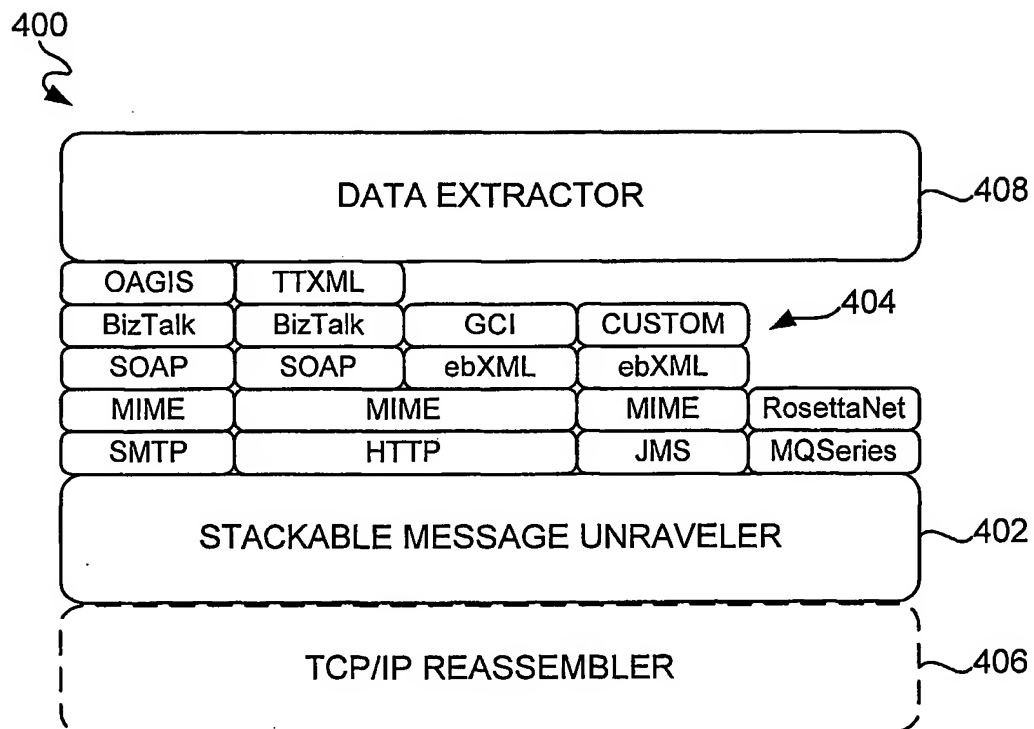


FIG. 3

**FIG. 4**




FIG. 5

6/7




FIG. 6




FIG. 7



## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US02/15163

<b>A. CLASSIFICATION OF SUBJECT MATTER</b>																				
IPC(7) : HO4L 9/00 US CL : 713/150,170,176,178 According to International Patent Classification (IPC) or to both national classification and IPC																				
<b>B. FIELDS SEARCHED</b>																				
Minimum documentation searched (classification system followed by classification symbols) U.S. : 713/150,170,176,178																				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched																				
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) Please See Extra Sheet.																				
<b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b>																				
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.																		
Y	US 6,012,098 A(BAYEH et al.) 04 January 2000, col.8,lines 13-29, col.9,lines 64-67,col.10,lines 1-15.	1-20																		
Y	US 5,235,642 A(WOBBER et al.) 10 August 1993, col.6,lines 23-39, col.7,lines 22-63.	1-20																		
A	US 5,533,021 A(BRANSTAD et al.) 02 July 1996, col.4,lines 41-58, col.5,lines 11-23.	1-20																		
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.																				
<table border="0"> <tr> <td>* Special categories of cited documents:</td> <td>"T"</td> <td>later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</td> </tr> <tr> <td>"A" document defining the general state of the art which is not considered to be of particular relevance</td> <td>"X"</td> <td>document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone</td> </tr> <tr> <td>"E" earlier document published on or after the international filing date</td> <td>"Y"</td> <td>document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</td> </tr> <tr> <td>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</td> <td>"&amp;"</td> <td>document member of the same patent family</td> </tr> <tr> <td>"O" document referring to an oral disclosure, use, exhibition or other means</td> <td></td> <td></td> </tr> <tr> <td>"P" document published prior to the international filing date but later than the priority date claimed</td> <td></td> <td></td> </tr> </table>			* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention	"A" document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone	"E" earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art	"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family	"O" document referring to an oral disclosure, use, exhibition or other means			"P" document published prior to the international filing date but later than the priority date claimed		
* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention																		
"A" document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone																		
"E" earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art																		
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family																		
"O" document referring to an oral disclosure, use, exhibition or other means																				
"P" document published prior to the international filing date but later than the priority date claimed																				
Date of the actual completion of the international search 14 JUNE 2002		Date of mailing of the international search report 12 JUL 2002																		
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230		Authorized officer GAIL HAYES <i>Peggy Hamod</i> Telephone No. (703) 305-0042																		

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/US02/15163

### B. FIELDS SEARCHED

Electronic data bases consulted (Name of data base and where practicable terms used):

STN,EAST

search terms: file,data,message,XML,HTML,stream,timestamp,digital  
signature,authenticate,verify,certificate,extract,encrypt,cipher,encipher